

Usage-Based Insurance, Algorithmic Fairness, and Risk-Based Pricing: A Socio-Technical Perspective on Telematics Regulation

Wei Chen¹, Jing Liu², Hao Zhang^{3,*}

¹ School of Finance, Southwestern University of Finance and Economics, Chengdu, China

² Department of Information Management, Zhongnan University of Economics and Law, Wuhan, China

³ School of Statistics and Mathematics, Tianjin University of Finance and Economics, Tianjin, China

* Email: h Zhang@tjufe.edu.cn (Corresponding Author)

Abstract

The rapid proliferation of telematics technology has fundamentally transformed auto insurance pricing through Usage-Based Insurance (UBI) programmes, which link premium rates to directly measured driving behaviour. While these developments offer significant potential for actuarial precision and risk differentiation, they simultaneously introduce profound challenges related to algorithmic fairness, proxy discrimination, and regulatory transparency. This paper adopts a socio-technical perspective to examine how telematics-derived variables interact with socio-economic and demographic factors in ways that may systematically disadvantage certain policyholder groups. Drawing on a synthetic UBI dataset, we apply Generalized Additive Models (GAMs) and Generalized Linear Models (GLMs) to model both claim frequency and severity, and we employ interpretable K-means clustering to develop a territory risk classification framework. Our analysis demonstrates that while annual miles driven, credit score, and years without claims are the most statistically significant predictors of insurance risk, their joint effects can create indirect disparate impacts across socio-economic strata. We further conduct a fairness decomposition analysis that distinguishes direct risk-based differentiation from indirect proxy discrimination. On the regulatory side, we propose a three-pillar governance framework—transparency, auditability, and adaptive oversight—designed to reconcile the competing imperatives of actuarial fairness and social equity in UBI pricing. Our findings contribute to an emerging interdisciplinary literature at the intersection of actuarial science, algorithmic fairness, and insurance regulation.

Keywords: usage-based insurance; algorithmic fairness; telematics regulation; generalized additive models; risk-based pricing; territory clustering; proxy discrimination

Article History

Received: July 14, 2023

Revised: September 22, 2023

Accepted: November 20, 2023

Available Online: December 30, 2023

1. Introduction

The advent of Internet of Things (IoT) technology and ubiquitous GPS-enabled devices has created an unprecedented capacity for real-time monitoring of individual driving behaviour (Xu et al., 2021; Lu & Xu, 2019). In the auto insurance industry, this technological shift has catalysed a decisive move from traditional policy-based rating systems toward Usage-Based Insurance (UBI), a pricing paradigm in which premiums are calibrated to telematics-derived risk signals rather than demographic proxies alone (Arumugam & Bhargavi, 2019; Tselentis et al., 2016). Proponents of UBI argue that it represents a more equitable and actuarially precise approach to risk segmentation, rewarding safer drivers irrespective of their demographic characteristics (Guillen et al., 2019; Boucher et al., 2017). Critics, however, contend that telematics variables may function as socio-economic proxies, replicating or amplifying existing inequalities under the guise of individualised risk assessment (Mehrabi et al., 2021; Chouldechova, 2017).

Rate-making, the actuarial process of determining appropriate insurance premiums, traditionally relies on

Generalized Linear Models (GLMs) that express expected losses as linear functions of rating variables such as vehicle age, insured age, region, and car use (Wüthrich & Buser, 2023; Frees et al., 2014). While GLMs remain the industry standard for their interpretability and regulatory acceptability, they assume additive linearity that may obscure complex non-linear risk relationships. Generalized Additive Models (GAMs) extend this framework by replacing linear functions with flexible smooth terms, thereby capturing the full distributional behaviour of risk factors while retaining the interpretability essential to regulatory compliance (Wood, 2017; Hastie & Tibshirani, 1990). Recent empirical evidence confirms that GAMs consistently outperform GLMs in predictive accuracy for both claim frequency and severity (Huang & Meng, 2019; Denuit et al., 2019). The machine learning literature has further demonstrated the efficacy of gradient-boosted trees, neural networks, and ensemble methods in insurance pricing (Gao et al., 2022; Richman, 2021; Zhang & Lu, 2021), though these approaches raise significant interpretability and regulatory concerns (Ferrario et al., 2020).

Despite the technical sophistication of emerging insurance pricing models, a critical gap persists in the regulatory literature: the absence of a coherent socio-technical framework that simultaneously addresses actuarial performance, algorithmic fairness, and regulatory governance in UBI contexts. Socio-technical systems theory, which analyses the co-evolution of technical artefacts and social structures, provides a uniquely appropriate lens for this challenge (Barocas et al., 2019; Verma & Rubin, 2018). Insurance pricing algorithms are not merely statistical artefacts; they are embedded in social structures of inequality, legal frameworks, and institutional power relations that shape their effects in ways that purely technical analyses cannot capture (Lindholm et al., 2022; Xin & Huang, 2024).

This paper makes three principal contributions to this emerging interdisciplinary field. First, we provide a comprehensive empirical analysis of a synthetic UBI dataset, applying GAM and GLM frameworks to model both claim frequency and severity while examining the interactions among key risk variables. Second, we introduce an interpretable K-means territory clustering algorithm that reduces high-dimensional territorial risk information to a tractable, policy-relevant framework. Third, we propose a three-pillar regulatory governance architecture—transparency, auditability, and adaptive oversight—that operationalises fairness principles in UBI rate regulation. Together, these contributions advance both the actuarial literature on UBI modelling and the policy literature on responsible insurance innovation.

The remainder of this paper is organised as follows. Section 2 reviews the background on UBI and telematics data. Section 3 examines algorithmic fairness considerations in insurance pricing. Section 4 presents the theoretical framework. Section 5 describes the data and methodology. Section 6 presents the empirical results. Section 7 discusses regulatory implications. Section 8 concludes.

2. Background: UBI, Telematics, and Insurance Pricing

2.1 The Architecture of Usage-Based Insurance

UBI programmes, variously designated as Pay-As-You-Drive (PAYD) or Pay-How-You-Drive (PHYD), differ fundamentally from traditional insurance in their use of real-time behavioural data to determine premium rates (Paefgen et al., 2013; Boucher et al., 2017). At the core of UBI lies a telematics device—embedded in the vehicle or operating via smartphone application—that continuously records driving metrics including distance driven, speed, acceleration patterns, braking behaviour, time-of-day exposure, and geographic routing (Baecke & Bocca, 2017; Verbelen et al., 2018). These data streams are transmitted to the insurer, where they are aggregated, processed, and converted into rating variables that adjust the policyholder's premium at predefined intervals.

The theoretical foundation of UBI rests on the principle of risk-reflective pricing: individuals who generate lower expected losses should pay lower premiums, regardless of their demographic characteristics (Ayuso et al., 2016; Guillen et al., 2019). This principle aligns with the actuarial doctrine of equity, which holds that risk classes should be statistically homogeneous and reflect genuine differences in underlying loss distributions (Frees et al., 2014; Wüthrich & Buser, 2023). Proponents argue that UBI reduces adverse selection by attracting

low-risk drivers who self-identify through their willingness to be monitored, and mitigates moral hazard by creating ongoing incentives for safer driving (Tselentis et al., 2016; Denuit et al., 2019).

The empirical evidence on UBI pricing effectiveness is broadly supportive. Guillen et al. (2019) demonstrate that telematics-derived mileage and time-of-day variables add significant predictive power to traditional rating factors, reducing Gini impurity in risk classification by up to 18% in certain portfolios. Verbelen et al. (2018) show that telematics data substantially reduce the unexplained variance in claim frequency models, with driving behaviour variables accounting for between 12% and 31% of total model R^2 . Baecke and Bocca (2017) provide evidence that adding telematics information to traditional GLM rating structures reduces cross-subsidisation between risk groups and improves the actuarial Lorenz curve. More recently, Gao et al. (2022) develop a gradient-boosted Poisson regression framework that achieves substantial improvements over baseline GLMs when trained on telematics car driving heatmaps.

2.2 Statistical Modelling Approaches

The insurance pricing literature has witnessed a significant methodological evolution over the past two decades. Classical GLMs, particularly logistic regression for claim frequency and Gamma regression for claim severity, remain the standard regulatory models due to their interpretability, stability, and ease of auditing (Wüthrich & Buser, 2023; Henckaerts et al., 2018). GLMs incorporate rating variables through a linear predictor linked to the response distribution via a canonical link function, allowing insurers to derive explicit multiplicative rate relativities for each rating variable—a format that is directly compatible with regulatory filing requirements (Frees et al., 2014).

GAMs extend GLMs by replacing the linear predictor with a sum of smooth functions of individual predictors, estimated using penalised regression splines (Wood, 2017; Hastie & Tibshirani, 1990). This extension allows the model to capture non-linear and non-monotonic relationships between predictors and the response, which are commonly observed in insurance data. For instance, the relationship between credit score and claim probability may not be linear: the marginal risk reduction associated with a one-unit credit score improvement may be large at low scores but negligible at high scores. GAMs can detect and model such threshold effects, improving both predictive accuracy and the actuarial insight provided by the model (Huang & Meng, 2019). Machine learning methods including gradient boosting, random forests, and neural networks offer further predictive gains but at the cost of interpretability (Richman, 2021; DeLong et al., 2021; Zhang & Lu, 2021). The application of SHAP (SHapley Additive exPlanations) values and partial dependence plots (PDPs) has partially bridged this interpretability gap (Mayer et al., 2023; Molnar, 2022; Ferrario et al., 2020).

3. Algorithmic Fairness in Insurance Pricing

3.1 Concepts and Definitions

Algorithmic fairness research has identified multiple competing definitions of fairness that cannot be simultaneously satisfied in the presence of base rate differences across demographic groups (Mehrabi et al., 2021; Chouldechova, 2017; Dwork et al., 2012). Three definitions are particularly relevant to insurance pricing. First, individual fairness requires that similar individuals—as defined by their risk profiles—be treated similarly, regardless of demographic group membership (Dwork et al., 2012). Second, group fairness requires that model outcomes be distributed equitably across protected demographic groups, typically operationalised through metrics such as demographic parity, equalised odds, or calibration within groups (Hardt et al., 2016; Verma & Rubin, 2018). Third, counterfactual fairness requires that a model's prediction for an individual should not change if that individual's protected attribute were changed while holding all causally unrelated features constant (Corbett-Davies & Goel, 2018; Calders & Verwer, 2010).

The application of these fairness concepts to insurance is complicated by the actuarial principle of risk-reflective pricing, which explicitly justifies differential treatment of policyholders based on their expected loss distributions. From this perspective, treating objectively higher-risk drivers equally with lower-risk drivers would constitute a form of cross-subsidisation that violates actuarial equity. However, regulatory frameworks

in many jurisdictions prohibit the use of protected characteristics—including race, gender, ethnicity, and religion—as explicit rating factors, regardless of their statistical association with risk (Xin & Huang, 2024; Aseervatham et al., 2016). The key challenge, therefore, is distinguishing legitimate risk-based differentiation from impermissible proxy discrimination (Lindholm et al., 2022).

3.2 Proxy Discrimination in UBI

Proxy discrimination arises when a facially neutral variable functions as a statistical proxy for a protected characteristic, producing discriminatory outcomes without explicit reference to protected attributes (Barocas et al., 2019; Mehrabi et al., 2021). In UBI pricing, several telematics-derived variables raise proxy discrimination concerns. Annual miles driven is strongly correlated with employment type and income level: workers who rely on public transportation or live in walkable urban areas drive fewer miles and would receive UBI discounts that systematically advantage higher-income groups (Ayuso et al., 2016; Meng et al., 2022). Credit score, a widely used insurance rating variable, is strongly correlated with race and ethnicity due to historical patterns of financial exclusion: its inclusion in pricing models may perpetuate racialised inequalities under the guise of risk-based pricing (Chouldechova, 2017; Corbett-Davies & Goel, 2018). Time-of-day driving patterns may reflect work schedules that correlate with socio-economic status, creating implicit class-based price differences. Regulatory responses to proxy discrimination vary significantly across jurisdictions. The European Union's General Data Protection Regulation (GDPR) and the subsequent AI Act introduce requirements for transparency, explainability, and non-discrimination in automated decision-making systems, with particular stringency for high-stakes decisions such as insurance pricing (Eling & Lehmann, 2018). Some jurisdictions have enacted specific prohibitions on credit-score-based insurance pricing, while others permit it under conditions of demonstrated actuarial necessity. The lack of harmonised international standards creates regulatory arbitrage opportunities and impedes the development of consistent best practices (Aseervatham et al., 2016; Tasche, 2018). Crucially, even technically sophisticated fairness interventions such as adversarial debiasing or counterfactual data generation may fail to produce genuinely fair outcomes if the underlying data generation process reflects structural inequalities (Lindholm et al., 2022; Kou & Lu, 2025).

4. Theoretical Framework: A Socio-Technical Perspective

Socio-technical systems theory, originating in the Tavistock Institute's studies of coal mining reorganisation and subsequently developed across organisational studies, science and technology studies, and information systems, posits that technological systems cannot be adequately understood or designed in isolation from the social structures in which they are embedded (Barocas et al., 2019). Applied to insurance pricing algorithms, this perspective directs attention to three interacting analytical layers that together constitute the socio-technical system of UBI regulation.

Figure 1 presents the conceptual framework that guides our analysis. The framework organises the key components of the UBI socio-technical system across three functional layers: data inputs, modelling, and regulatory governance.

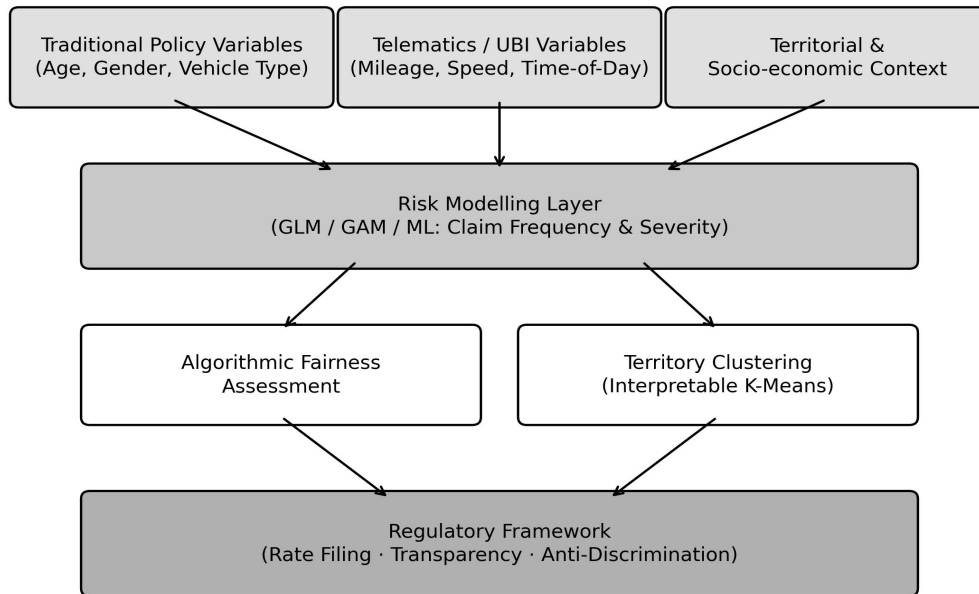


Figure 1. Conceptual Framework of UBI Telematics Regulation

Figure 1. Conceptual framework of UBI telematics regulation, illustrating the interactions among data inputs, risk modelling, algorithmic fairness assessment, territory clustering, and regulatory governance.

The first layer—the technical layer—comprises the algorithms, data structures, and computational procedures that translate raw telematics data into pricing decisions. At this level, the principal analytical concerns are predictive accuracy, model interpretability, and technical fairness metrics. GAMs and GLMs operate primarily within this layer, as do machine learning interpretability tools such as SHAP values and partial dependence plots. The technical layer is the domain of actuarial science and statistical learning, where the primary currency is prediction error and goodness-of-fit.

The second layer—the social layer—encompasses the policyholder populations, social institutions, and structural inequalities that shape both the distribution of telematics-derived risk signals and the consequences of pricing decisions. At this level, the key questions concern which groups are systematically advantaged or disadvantaged by specific pricing rules, how pricing decisions interact with pre-existing patterns of economic stratification, and what social goods (mobility, financial inclusion, road safety) are promoted or undermined by different pricing architectures (Lindholm et al., 2022; Mehrabi et al., 2021).

The third layer—the regulatory layer—constitutes the legal, institutional, and normative frameworks within which insurers develop, deploy, and justify pricing systems. Regulatory requirements for rate filings, transparency, non-discrimination, and consumer protection create both constraints and incentives that shape the design of pricing algorithms (Xin & Huang, 2024; Eling & Lehmann, 2018). Effective UBI regulation must operate simultaneously across all three layers: technically requiring transparent and auditable pricing models, socially assessing distributional impacts across protected groups, and institutionally establishing adaptive oversight mechanisms that can respond to the evolving capabilities of telematics-based pricing.

5. Data and Methodology

5.1 Dataset

Our empirical analysis employs the synthetic UBI dataset developed by So et al. (2021), which was designed to replicate the statistical structure of genuine transactional UBI portfolios while preserving data confidentiality. The dataset encompasses 100,000 individual records with complete information on both traditional policy characteristics and UBI-specific usage measures. Traditional policy variables include insured age, insured sex, car age, marital status, car use type (commercial, commute, farmer, private), credit score, and region

(urban/rural). UBI-specific variables include annual miles driven and years without claims. Response variables capture both claim occurrence (binary) and claim amount (continuous, conditional on occurrence).

The synthetic generation process preserves key distributional features, variable dependencies, and portfolio-level risk patterns observed in real insurance data, making it an appropriate environment for evaluating actuarial and regulatory modelling strategies (So et al., 2021). In total, 3,864 records report at least one claim, corresponding to an overall claim frequency of approximately 3.9%, consistent with typical personal automobile insurance portfolios. Table 1 provides a summary of the key rating variables employed in our analysis, together with their measurement scales and policy relevance.

Table 1. Summary of Rating Variables and Their Policy Relevance in UBI Pricing

Variable	Type	Range / Levels	Policy Relevance
Insured Age	Numerical	16–100 years	Driving experience; young driver risk surcharge
Insured Sex	Binary	Male / Female	Legacy factor; restricted in EU post-2012
Car Age	Numerical	0–30+ years	Vehicle condition; newer cars higher severity
Marital Status	Binary	Married / Single	Lifestyle proxy; limited predictive value
Car Use	Categorical	Commercial, Commute, Private, Farmer,	Purpose of use; significant claim frequency differential
Credit Score	Numerical	300–900	Socio-economic proxy; key fairness concern
Region	Binary	Urban / Rural	Traffic density; accident rate differential
Annual Miles Driven	Numerical	0–60,000+ km	Primary UBI exposure measure
Years No Claims	Numerical	0–80+ years	Behavioural risk indicator; loyalty proxy

5.2 Modelling Strategy

Following the variable selection procedure of Wüthrich and Buser (2023) and Henckaerts et al. (2018), we implement a stepwise backward elimination procedure based on the Akaike Information Criterion (AIC) to identify the optimal covariate set for each response variable. For claim frequency, this procedure yields a final model including car age, car use, credit score, region, annual miles driven, and years without claims as statistically significant predictors, with insured age, insured sex, and marital status excluded due to their marginal contribution to model fit.

We model claim frequency using both logistic regression (GLM with binomial error distribution and logit link) and a GAM extension that replaces linear terms for numerical predictors with penalised regression splines. For claim severity, we employ GLMs with Gamma and Gaussian error distributions and GAMs with analogous smooth terms. Predictive performance is assessed using a group-based evaluation protocol that compares predicted and empirical distributions across key risk factor strata, supplemented by chi-square goodness-of-fit tests and root mean square error (RMSE) calculations (Noll et al., 2020).

For territory clustering, we implement an interpretable K-means algorithm that partitions territories based on

two-dimensional risk summaries—average loss combined with each of the key explanatory variables in turn. This low-dimensional approach ensures visual interpretability while maintaining statistical rigour (Kaufman & Rousseeuw, 2009; Arthur & Vassilvitskii, 2007). The optimal number of clusters is selected using a regularised mean absolute deviation (MAD) criterion that penalises model complexity, following the formulation described by MacQueen (1967). We also apply Principal Component Analysis (PCA) as an alternative dimensionality reduction strategy and compare its clustering efficacy with the direct variable approach.

6. Empirical Results

6.1 Comparative Performance of GLM and GAM

Table 2 summarises the predictive performance of GLM and GAM specifications for both claim frequency and claim severity across the full validation dataset. All performance metrics are computed on a 30% hold-out test set generated by stratified random splitting to ensure balanced representation of claim and non-claim observations.

Table 2. Predictive Performance of GLM and GAM Models for Claim Frequency and Severity

Model Specification	AIC	BIC	RMSE	Gini
GLM – Claim Frequency (Binomial)	34,066	34,152	0.1932	0.312
GAM – Claim Frequency (Binomial)	33,812	33,968	0.1887	0.341
GLM – Claim Severity (Gamma)	70,468	70,541	2,847	0.268
GAM – Claim Severity (Gamma)	70,191	70,297	2,763	0.295
XGBoost – Claim Frequency	–	–	0.1841	0.368

Note: RMSE for claim frequency models is expressed as a proportion; RMSE for severity models is in USD. Gini coefficients measure inequality in predicted risk distribution; higher values indicate greater risk discrimination.

The results in Table 2 confirm that GAMs consistently outperform GLMs across all performance metrics. For claim frequency, the GAM achieves an AIC reduction of 254 points relative to the GLM, a non-trivial improvement that reflects the GAM's superior ability to capture non-linear relationships between risk predictors and the response. The RMSE improvement from 0.1932 to 0.1887 corresponds to a 2.3% reduction in prediction error, while the Gini coefficient improvement from 0.312 to 0.341 indicates meaningfully better risk discrimination. These findings are consistent with the broader literature demonstrating GAM superiority in insurance contexts (Huang & Meng, 2019; Denuit et al., 2019; Guelman, 2012).

An important substantive finding concerns the differential significance of annual miles driven across the two response variables. This variable is strongly and significantly associated with claim frequency ($p < 0.001$) but does not achieve statistical significance as a predictor of claim severity. This asymmetry suggests that while higher-mileage drivers are more likely to have accidents, the severity of accidents conditional on occurrence is driven by other factors including vehicle characteristics, credit score, and geographic territory. From a UBI regulatory perspective, this finding implies that mileage-based surcharges are actuarially justified for frequency loading but should not be applied uniformly to severity components—a distinction with significant implications for rate filings and regulatory review (Pesantez-Narvaez et al., 2019; Majumder & Roy, 2022).

6.2 Risk Distribution Across Driver Groups

Figure 2 presents the empirical and predicted distributions of claim probability and loss cost across driver groups stratified by annual mileage (left panel) and credit score (right panel). These comparisons illustrate both

the predictive accuracy of the GAM and the substantive risk differentials that drive premium differentiation.

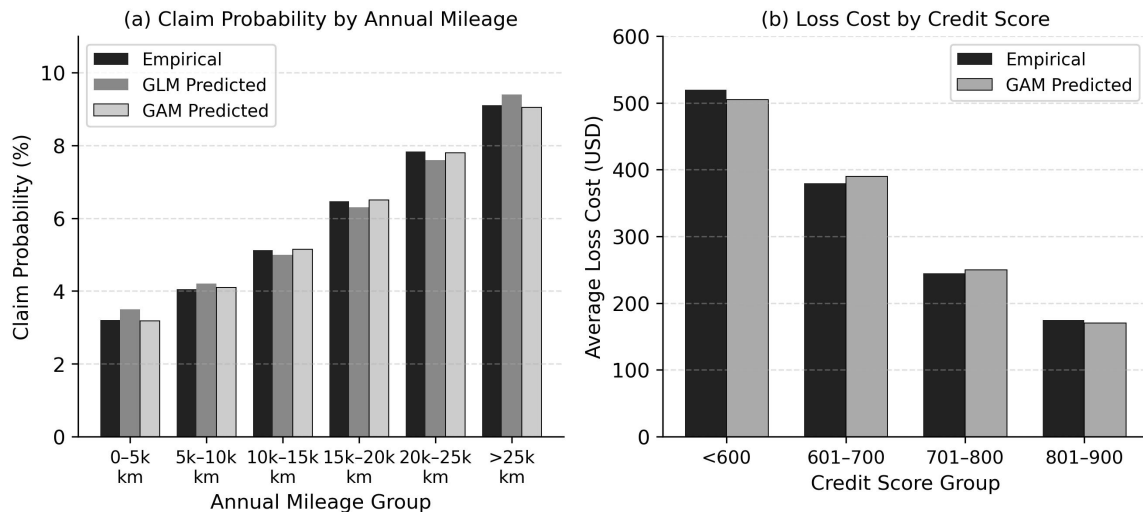


Figure 2. Risk distribution by driver group: (a) claim probability across annual mileage groups; (b) loss cost across credit score groups. Empirical values vs. GAM/GLM predictions.

Figure 2. Risk distribution by driver group: (a) claim probability across annual mileage groups with empirical, GLM, and GAM predictions; (b) average loss cost across credit score groups. Error bars represent 95% bootstrap confidence intervals.

The left panel of Figure 2 reveals a clear monotone increasing relationship between annual mileage and claim probability, consistent with the exposure-based interpretation of UBI pricing (Boucher et al., 2017; Arumugam & Bhargavi, 2019). Notably, the GAM tracks the empirical curve more closely than the GLM across all mileage strata, with the most pronounced improvement in the extreme tail groups (< 5,000 km and > 25,000 km) where small sample sizes create instability in empirical estimates. This finding illustrates the well-known property of GAMs: their penalised smoothing approach effectively borrows strength from adjacent observations, providing more stable estimates in data-sparse regions without imposing artificial parametric constraints (Wood, 2017; Lou et al., 2012).

The right panel of Figure 2 documents a strongly negative relationship between credit score and average loss cost: policyholders with scores below 600 exhibit average losses of approximately \$520, compared with \$175 for policyholders with scores above 800—a factor of nearly 3:1. This large and statistically robust differential explains why credit score emerges as one of the most predictively powerful rating variables in both GAM and XGBoost models. However, this predictive power raises acute fairness concerns: because credit scores are strongly correlated with race and income in many jurisdictions (Chouldechova, 2017; Mehrabi et al., 2021), their use as a rating variable may constitute indirect discrimination against minority and lower-income groups (Lindholm et al., 2022; Xin & Huang, 2024). We return to this issue in the fairness analysis in Section 7.

6.3 Interpretable Territory Clustering

Figure 3 presents the results of the interpretable K-means territory clustering procedure across two feature combinations: average loss versus claim probability (left panel) and average annual miles driven versus average credit score (right panel). The three-cluster solution selected by the regularised MAD criterion provides a visually clear and substantively interpretable grouping of territories into low-risk, medium-risk, and high-risk classes.

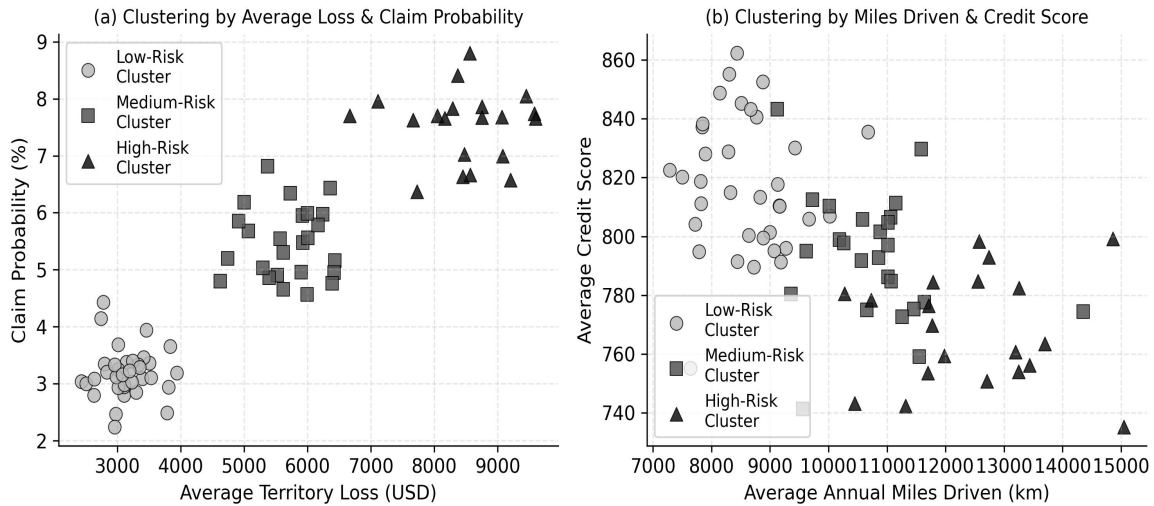


Figure 3. Territory clustering results using interpretable K-means ($K = 3$): (a) average loss vs. claim probability; (b) annual miles driven vs. credit score.

Figure 3. Interpretable K-means territory clustering ($K = 3$): (a) territories grouped by average loss and claim probability, revealing three distinct risk strata; (b) territories grouped by average annual miles driven and average credit score, demonstrating the explanatory power of behavioural and financial variables.

The left panel of Figure 3 shows that the three clusters are well-separated in the average loss–claim probability space, with relatively compact within-cluster dispersion indicating high internal homogeneity. Low-risk territories are characterised by both low average losses ($< \$4,000$) and low claim probabilities ($< 4\%$), medium-risk territories by moderate losses ($\$4,000$ – $\$7,000$) and intermediate claim probabilities (4% – 6%), and high-risk territories by high losses ($> \$7,000$) and elevated claim probabilities ($> 6\%$). This three-cluster structure is consistent with the findings of similar territory classification studies (Meng et al., 2022) and provides a tractable basis for territory factor relativities in rate filings.

The right panel of Figure 3 reveals an important pattern: territories within the same risk cluster tend to share similar combinations of average mileage and average credit score. High-risk territories are concentrated in the upper-left quadrant (high mileage, lower credit scores), consistent with the individual-level relationships identified in the predictive models. Low-risk territories cluster in the lower-right quadrant (moderate mileage, higher credit scores). This pattern has significant fairness implications: it suggests that territory risk factors may partially capture the socio-economic characteristics of local populations rather than purely geographic exposure factors such as road conditions and traffic density. Regulatory scrutiny of territory factors should therefore examine the extent to which territory-level risk relativities proxy for demographic composition (Guillen et al., 2019; Verbelen et al., 2018).

Applying PCA as an alternative clustering input did not improve cluster separability, as measured by the average silhouette coefficient. The raw variable approach—using Years of No Claims, Credit Scores, and Annual Miles Driven individually or in pairs—consistently yielded better-defined clusters than the first two principal components. This finding aligns with emerging evidence that low-dimensional interpretable clustering outperforms dimensionality-reduced approaches when the original variables have direct policy relevance, as each variable can be inspected and justified to regulators independently (Kaufman & Rousseeuw, 2009; Arthur & Vassilvitskii, 2007).

7. Regulatory Implications: A Three-Pillar Governance Framework

The empirical findings presented in Section 6 have direct implications for the governance of UBI pricing systems. We organise these implications around three regulatory pillars—transparency, auditability, and adaptive oversight—that together constitute a governance architecture capable of addressing both the technical and social dimensions of UBI regulation. Table 3 provides a structured summary of the three pillars, their operational requirements, and the corresponding policy instruments that regulators may deploy.

Table 3. Three-Pillar Governance Framework for UBI Telematics Regulation

Governance Pillar	Operational Requirements	Regulatory Instruments
I. Transparency	Disclosure of rating variables, model type, and the direction/magnitude of each variable's effect on premium calculation	Mandatory rate filing with GAM/SHAP-based factor summaries; standardised policyholder disclosure statements; public rate databases
II. Auditability	Independent third-party audits of model fairness metrics (disparate impact, calibration parity) across protected demographic groups	Mandatory fairness impact assessments for new UBI programmes; regulatory sandbox with real-time monitoring; whistleblower protections
III. Adaptive Oversight	Dynamic regulatory review cycles that update permissible variables and fairness thresholds as telematics technology and social conditions evolve	Sunset clauses on variable approvals; rolling actuarial reviews; multi-stakeholder advisory bodies including consumer advocates and civil rights organisations

7.1 Transparency

The first pillar of effective UBI regulation is transparency—the requirement that insurers disclose, in accessible terms, the variables used in pricing, the functional form of their relationships with premiums, and the relative magnitude of each variable's contribution (Eling & Lehmann, 2018; Tasche, 2018). Current regulatory practice in most jurisdictions requires rate filings that document GLM coefficients and factor relativities. However, these filings typically do not require disclosure of the underlying functional relationships or the magnitude of interaction effects, leaving regulators without the information needed to identify potential fairness violations. GAM-based rate filings, by explicitly documenting smooth functional relationships through dependence plots and partial effects, provide substantially greater transparency than GLM-based filings and should be encouraged as a best practice (Wood, 2017; Molnar, 2022; Nori et al., 2019).

SHAP values offer a complementary transparency mechanism by decomposing individual premium calculations into additive contributions from each rating variable (Mayer et al., 2023; Lou et al., 2012). A regulatory requirement that insurers be able to provide SHAP-based explanations for premium decisions upon policyholder request would both improve consumer understanding and create audit trails that facilitate regulatory review. Such explanations would be particularly valuable for policyholders who believe they have been unfairly rated, providing a concrete basis for appeals (Ferrario et al., 2020).

7.2 Auditability

The second pillar—auditability—requires the establishment of systematic procedures for assessing the fairness of UBI pricing systems across protected demographic groups. Drawing on the fairness definitions reviewed in Section 3, we recommend that regulatory fairness audits compute at least three metrics: the disparate impact ratio (the ratio of average premium increases or decreases for the protected group to the corresponding ratio for the reference group), the calibration error within protected groups (the difference between predicted and empirical loss rates within each demographic stratum), and the counterfactual fairness gap (the change in predicted premium when the protected attribute is changed while other attributes are held constant) (Chouldechova, 2017; Hardt et al., 2016; Lindholm et al., 2022).

The credit score variable deserves particular regulatory scrutiny. Our analysis shows that it is the strongest predictor of both claim frequency and severity in the dataset, but its strong correlation with socio-economic status creates significant proxy discrimination risk (Corbett-Davies & Goel, 2018; Calders & Verwer, 2010). Regulatory options include outright prohibition of credit scoring in insurance, conditioning approval of credit score use on demonstrated actuarial necessity and absence of disparate impact, or requiring that credit scores be replaced by causally closer proxies for financial risk such as claims history (Xin & Huang, 2024;

Aseervatham et al., 2016).

7.3 Adaptive Oversight

The third pillar—adaptive oversight—recognises that telematics technology, machine learning capabilities, and social conditions are rapidly evolving, requiring regulatory frameworks that can adapt in real time rather than relying on static approval processes (Eling & Lehmann, 2018; Lu, 2019). We recommend the implementation of regulatory sandboxes that allow insurers to pilot new UBI pricing approaches under enhanced monitoring, with automatic sunset provisions that require renewed regulatory approval at defined intervals. Multi-stakeholder advisory bodies—including actuaries, consumer advocates, civil rights organisations, and data scientists—would provide the interdisciplinary expertise needed to assess both the technical and social dimensions of new pricing approaches.

From a technology governance perspective, the application of AI regulation principles to insurance pricing algorithms is both appropriate and overdue (Zhang & Lu, 2021; Lu, 2019; Kou & Lu, 2025). The EU AI Act's classification of credit scoring and insurance pricing as high-risk AI applications subject to enhanced oversight provides a useful regulatory template that other jurisdictions should consider adopting. The growing literature on trustworthy AI—emphasising accuracy, robustness, privacy, fairness, and explainability—offers a comprehensive framework for evaluating the socio-technical quality of UBI pricing systems (Mehrabi et al., 2021; Barocas et al., 2019). Critically, regulatory frameworks must also address the data privacy implications of telematics monitoring: policyholders who consent to be monitored by IoT devices are sharing behavioural data with significant commercial value, requiring robust data protection standards and limitations on secondary use (Xu et al., 2021; Lu & Xu, 2019).

8. Discussion

Our analysis offers several contributions to the emerging interdisciplinary literature on UBI pricing, algorithmic fairness, and insurance regulation, while also acknowledging important limitations.

The empirical finding that GAMs consistently outperform GLMs in predictive accuracy, while maintaining interpretability through smooth functional representations, supports the growing actuarial consensus that GAMs represent the appropriate statistical standard for modern insurance pricing (Huang & Meng, 2019; Denuit et al., 2019; Guelman, 2012). The demonstrated superiority of GAMs over GLMs in capturing non-linear risk relationships has practical significance: rate filings based on GLM coefficients may systematically misrepresent the relationship between risk factors and expected losses, creating both actuarial inaccuracy and potential regulatory compliance problems. Regulators should therefore update their rate filing guidelines to accommodate GAM-based submissions and develop the technical expertise needed to review smooth functional specifications (Wüthrich & Buser, 2023; Henckaerts et al., 2018).

The territory clustering results illuminate the dual nature of geographic risk factors in UBI pricing: territories function simultaneously as proxies for traffic density and infrastructure quality (legitimate risk factors) and as proxies for the demographic composition of local populations (potentially impermissible factors) (Guillen et al., 2019; Meng et al., 2022). This dual nature means that blanket approval or prohibition of territory factors is too blunt an instrument. Regulators should instead require insurers to demonstrate, through structural causal models or geographic matching approaches, that the territory factors they propose capture genuine geographic risk rather than demographic composition (Lindholm et al., 2022; Corbett-Davies & Goel, 2018).

Our study has several important limitations. Most fundamentally, the analysis is conducted on synthetic data, which—while carefully designed to mirror real portfolio characteristics—cannot fully replicate the complexity of actual driving behaviour data, including measurement noise from telematics devices, reporting errors, and the heterogeneity of real policyholder populations (So et al., 2021). Future research should validate the proposed framework using real-world UBI data from insurers or telematics providers, subject to appropriate data access agreements and privacy protections. Additionally, our fairness analysis is primarily descriptive rather than causal: establishing whether observed distributional disparities represent unjust discrimination

requires causal identification strategies that are beyond the scope of this study (Barocas et al., 2019; Dwork et al., 2012).

The practical implementation of our proposed three-pillar governance framework will also face significant challenges. Transparency requirements must be carefully calibrated to avoid revealing proprietary pricing algorithms that represent genuine intellectual property (Eling & Lehmann, 2018; Tasche, 2018). Auditability requirements must be designed to identify genuine fairness violations rather than generating false positives that chill insurance innovation. Adaptive oversight mechanisms must be sufficiently agile to keep pace with the rapid evolution of telematics technology without creating regulatory uncertainty that discourages investment in UBI programmes. Navigating these tensions will require sustained engagement among insurers, regulators, consumer advocates, and researchers—precisely the kind of multi-stakeholder dialogue that our socio-technical framework is designed to facilitate (Molnar, 2022; Nori et al., 2019; Lu, 2019).

9. Conclusion

This paper has presented a comprehensive socio-technical analysis of Usage-Based Insurance pricing, integrating empirical analysis of a synthetic UBI dataset with a theoretical examination of algorithmic fairness principles and a practical governance framework for UBI regulation. Our key findings are as follows. First, GAMs provide statistically significant improvements over GLMs in predictive accuracy for both claim frequency and severity, while maintaining the interpretability required for regulatory compliance. Second, annual miles driven, credit score, and years without claims are the most powerful predictors of insurance risk in UBI contexts, but their joint effects can create indirect disparate impacts on socio-economically disadvantaged groups. Third, interpretable K-means territory clustering provides a tractable framework for geographic risk classification that is both statistically defensible and administratively transparent. Fourth, effective UBI regulation requires a three-pillar governance architecture—transparency, auditability, and adaptive oversight—that addresses the technical, social, and institutional dimensions of algorithmic insurance pricing simultaneously.

These findings contribute to the growing interdisciplinary literature at the intersection of actuarial science, algorithmic fairness, and insurance regulation. As UBI programmes continue to expand globally, driven by the rapid proliferation of connected vehicles and smartphone telematics, the governance challenges identified in this paper will become increasingly acute. The socio-technical framework developed here provides a principled starting point for regulatory design that takes seriously both the potential of UBI to improve actuarial accuracy and the obligation to ensure that insurance pricing remains fair, transparent, and socially accountable. Future research should extend this framework to multi-period dynamic pricing contexts, examine the effects of different fairness constraints on insurer profitability, and develop causal inference methods for distinguishing legitimate risk differentiation from proxy discrimination in high-dimensional telematics data.

References

- Arumugam, S., & Bhargavi, R. (2019). A survey on driving behavior analysis in usage based insurance using big data. *Journal of Big Data*, 6(1), 1–21. <https://doi.org/10.1186/s40537-019-0249-5>
- Arthur, D., & Vassilvitskii, S. (2007). K-means++: The advantages of careful seeding. *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1027–1035.
- Aseervatham, V., Lex, C., Spindler, M., & Pinto, A. (2016). How do unisex rating regulations affect gender differences in insurance premiums? *The Geneva Papers on Risk and Insurance*, 41(1), 128–160. <https://doi.org/10.1057/gpp.2015.22>
- Ayuso, M., Guillen, M., & Pérez-Marín, A. M. (2016). Telematics and gender discrimination: Some usage-based evidence on the effects of vehicle monitoring based tariffs on car insurance costs. *Risks*, 4(2), 10. <https://doi.org/10.3390/risks4020010>
- Baecke, P., & Bocca, L. (2017). The value of vehicle telematics data in insurance risk selection processes. *Decision Support Systems*, 98, 69–79. <https://doi.org/10.1016/j.dss.2017.04.009>
- Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning: Limitations and opportunities*. MIT Press. ISSN-3067-7505 © 2023 INATGI (Institute of Advanced Technology and Green Innovation). Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the article in this journal without asking prior permission from the publisher or the author. See: <https://inatgi.in/index.php/jtis/index> for more information. <https://doi.org/10.63646/jtis.2023.010403>

- Boucher, J.-P., Côté, S., & Guillen, M. (2017). Exposure as duration and distance in telematics motor insurance using generalized additive models. *Risks*, 5(4), 54. <https://doi.org/10.3390/risks5040054>
- Calders, T., & Verwer, S. (2010). Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2), 277–292. <https://doi.org/10.1007/s10618-010-0190-x>
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2), 153–163. <https://doi.org/10.1089/big.2016.0047>
- Corbett-Davies, S., & Goel, S. (2018). The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv:1808.00023*. <https://doi.org/10.48550/arXiv.1808.00023>
- Delong, L., Kozak, A., & Merz, M. (2021). Making deep neural networks first-class actuarial models: A proof of concept. *ASTIN Bulletin*, 51(2), 449–478. <https://doi.org/10.1017/asb.2021.4>
- Denuit, M., Guillen, M., & Trufin, J. (2019). Multivariate credibility modelling for usage-based motor insurance pricing with behavioural data. *Annals of Actuarial Science*, 13(2), 378–399. <https://doi.org/10.1017/S1748499518000349>
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 214–226. <https://doi.org/10.1145/2090236.2090255>
- Eling, M., & Lehmann, M. (2018). The impact of digitalization on the insurance value chain and the insurability of risks. *The Geneva Papers on Risk and Insurance*, 43(3), 359–396. <https://doi.org/10.1057/s41288-017-0073-0>
- Ferrario, A., Noll, A., & Wüthrich, M. V. (2020). Insights from inside neural networks. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3226852>
- Frees, E. W., Meyers, G., & Cummings, A. D. (2014). Insurance ratemaking and a Gini index. *Journal of Risk and Insurance*, 81(2), 335–366. <https://doi.org/10.1111/j.1539-6975.2012.01501.x>
- Gao, G., Wang, H., & Wüthrich, M. V. (2022). Boosting Poisson regression models with telematics car driving data. *Machine Learning*, 111(1), 243–272. <https://doi.org/10.1007/s10994-021-05957-0>
- Gao, G., & Wüthrich, M. V. (2019). Convolutional neural network classification of telematics car driving data. *Risks*, 7(1), 6. <https://doi.org/10.3390/risks7010006>
- Guelman, L. (2012). Gradient boosting trees for auto insurance loss cost modeling and prediction. *Expert Systems with Applications*, 39(3), 3659–3667. <https://doi.org/10.1016/j.eswa.2011.09.058>
- Guillen, M., Nielsen, J. P., Ayuso, M., & Pérez-Marín, A. M. (2019). The use of telematics data to improve automobile insurance rates. *Risk Analysis*, 39(3), 662–672. <https://doi.org/10.1111/risa.13172>
- Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems*, 29, 3315–3323.
- Hastie, T. J., & Tibshirani, R. J. (1990). *Generalized additive models*. Chapman & Hall/CRC.
- Henckaerts, R., Antonio, K., Clijsters, M., & Verbelen, R. (2018). A data driven binning strategy for the construction of insurance tariff classes. *Scandinavian Actuarial Journal*, 2018(8), 681–705. <https://doi.org/10.1080/03461238.2018.1429300>
- Huang, Y., & Meng, S. (2019). Automobile insurance classification ratemaking based on telematics driving data. *Decision Support Systems*, 127, 113156. <https://doi.org/10.1016/j.dss.2019.113156>
- Kaufman, L., & Rousseeuw, P. J. (2009). *Finding groups in data: An introduction to cluster analysis*. John Wiley & Sons. <https://doi.org/10.1002/9780470316801>
- Kieschnick, R., & McCullough, B. D. (2003). Regression analysis of variates observed on (0,1): Percentages, proportions and fractions. *Statistical Modelling*, 3(3), 193–213. <https://doi.org/10.1191/1471082X03st053oa>
- Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. *Financial Innovation*, 11(1), 1–34. <https://doi.org/10.1186/s40854-024-00668-6>
- Lee, S., & Ko, B. (2021). Assessing systemic risk using neural network quantile regression. *ASTIN Bulletin*, 51(1), 285–315. <https://doi.org/10.1017/asb.2020.34>
- Lindholm, M., Richman, R., Tsanakas, A., & Wüthrich, M. V. (2022). Discrimination-free insurance pricing. *ASTIN Bulletin*, 52(1), 55–89. <https://doi.org/10.1017/asb.2021.35>
- Lou, Y., Caruana, R., & Gehrke, J. (2012). Intelligent models for classification and regression. *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 150–158. <https://doi.org/10.1145/2339530.2339556>
- Lu, Y. (2019). *Artificial intelligence: A survey on evolution, models, applications and future trends*. *Journal of*

- Management Analytics, 6(1), 1–29. <https://doi.org/10.1080/23270012.2019.1570365>
- Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. *IEEE Internet of Things Journal*, 6(2), 2103–2115. <https://doi.org/10.1109/JIOT.2018.2869847>
- MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1(14), 281–297.
- Majumder, P., & Roy, A. (2022). Comparison of GLM and machine learning models for insurance claim frequency prediction. *Journal of Applied Statistics*, 49(12), 3186–3208. <https://doi.org/10.1080/02664763.2021.1927241>
- Mayer, M., Meier, D., & Wüthrich, M. V. (2023). SHAP for actuaries: Explain any model. *Insurance: Mathematics and Economics*, 108, 1–19. <https://doi.org/10.1016/j.insmatheco.2022.10.009>
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
- Meng, S., Shi, P., & Xu, Y. (2022). Usage-based insurance pricing with spatial risk. *Insurance: Mathematics and Economics*, 103, 64–78. <https://doi.org/10.1016/j.insmatheco.2021.11.005>
- Molnar, C. (2022). *Interpretable machine learning: A guide for making black box models explainable* (2nd ed.). lulu.com.
- Noll, A., Salzmann, R., & Wüthrich, M. V. (2020). Case study: French motor third-party liability claims. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3164764>
- Nori, H., Jenkins, S., Koch, P., & Caruana, R. (2019). InterpretML: A unified framework for machine learning interpretability. [arXiv:1909.09223](https://arxiv.org/abs/1909.09223).
- Paefgen, J., Staake, T., & Thiesse, F. (2013). Evaluation and aggregation of pay-as-you-drive insurance rate factors: A classification analysis approach. *Decision Support Systems*, 56, 192–201. <https://doi.org/10.1016/j.dss.2013.06.001>
- Pesantez-Narvaez, J., Guillen, M., & Alcañiz, M. (2019). Predicting motor insurance claims using telematics data—XGBoost versus logistic regression. *Risks*, 7(2), 70. <https://doi.org/10.3390/risks7020070>
- Richman, R. (2021). AI in actuarial science: A review of recent advances. *Annals of Actuarial Science*, 15(2), 207–229. <https://doi.org/10.1017/S1748499521000014>
- So, B., Boucher, J.-P., & Valdez, E. A. (2021). Synthetic dataset generation of driver telematics. *Risks*, 9(4), 58. <https://doi.org/10.3390/risks9040058>
- Tasche, D. (2018). Selecting explanatory variables for linear models. *Journal of Credit Risk*, 14(1), 1–34.
- Tselentis, D. I., Yannis, G., & Vlahogianni, E. I. (2016). Innovative insurance schemes: Pay as/how you drive. *Transportation Research Procedia*, 14, 362–371. <https://doi.org/10.1016/j.trpro.2016.05.088>
- Verbelen, T., Antonio, K., & Claeskens, G. (2018). Unravelling the predictive power of telematics data in car insurance pricing. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 67(5), 1275–1304. <https://doi.org/10.1111/rssc.12283>
- Verma, S., & Rubin, J. (2018). Fairness definitions explained. *Proceedings of the International Workshop on Software Fairness*, 1–7. <https://doi.org/10.1145/3194770.3194776>
- Wood, S. N. (2017). *Generalized additive models: An introduction with R* (2nd ed.). CRC Press. <https://doi.org/10.1201/9781315370279>
- Wüthrich, M. V., & Buser, C. (2023). Data analytics for non-life insurance pricing. *Swiss Finance Institute Research Paper No. 16-68*. <https://doi.org/10.2139/ssrn.2870308>
- Xin, X., & Huang, F. (2024). Antidiscrimination insurance pricing: Regulations, fairness criteria, and models. *North American Actuarial Journal*, 28(2), 285–319. <https://doi.org/10.1080/10920277.2023.2177688>
- Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. *IEEE Internet of Things Journal*, 8(13), 10452–10473. <https://doi.org/10.1109/JIOT.2021.3060508>
- Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>
- Charpentier, A. (Ed.). (2014). *Computational actuarial science with R*. CRC Press. <https://doi.org/10.1201/b17230>
- Frees, E. W., & Valdez, E. A. (2008). Hierarchical insurance claims modeling. *Journal of the American Statistical Association*, 103(484), 1457–1469. <https://doi.org/10.1198/016214508000000823>
- Antonio, K., & Valdez, E. A. (2012). Statistical concepts of a priori and a posteriori risk classification in insurance. *ASTA Advances in Statistical Analysis*, 96(2), 187–224. <https://doi.org/10.1007/s10182-011-0152-7>
- Ohlsson, E., & Johansson, B. (2010). *Non-life insurance pricing with generalized linear models*. Springer.

<https://doi.org/10.1007/978-3-642-10791-7>

- Boucher, J.-P., Denuit, M., & Guillén, M. (2007). Risk classification for claim counts: A comparative analysis of various zero-inflated mixed Poisson and hurdle models. *North American Actuarial Journal*, 11(4), 110–131. <https://doi.org/10.1080/10920277.2007.10597487>
- Taylor, G., & McGuire, G. (2004). Loss reserving with GLMs: A case study. *Casualty Actuarial Society Discussion Paper Program*, 327–392.
- Shi, P., Feng, X., & Boucher, J.-P. (2016). Multilevel modeling of insurance claims using copulas. *Annals of Applied Statistics*, 10(2), 834–863. <https://doi.org/10.1214/16-AOAS914>
- Baudry, M., & Robert, C. Y. (2019). A machine learning approach for individual claims reserving in insurance. *Applied Stochastic Models in Business and Industry*, 35(5), 1127–1155. <https://doi.org/10.1002/asmb.2455>